> Adventures with illumos

Peter Tribble

Theoretical Astrophysicist

Sysadmin

(DBA)

Technology Tinkerer

> Introduction

- Long-time systems administrator
- Many years pointing out bugs in Solaris
- Invited onto beta programs
- Then the OpenSolaris project
- Voted onto OpenSolaris Governing Board
- Along came Oracle...
- illumos emerged from the ashes

> illumos key differentiators

- ZFS – reliable and easy to manage
- Dtrace – extreme observability
- Zones – lightweight virtualization
- Standards – pretty strict
- Compatibility – decades of heritage
- “Solarishness”
- No systemd ;-)

> Diverse distributions

- OpenIndiana – OpenSolaris
- OmniOS – server focus
- SmartOS – Joyent's cloud
- Delphix/Nexenta/+ – storage focus
- DilOS, osdyson – debian
- Tribblix – one of the small fry

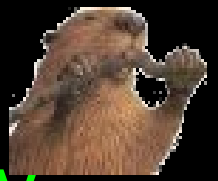

XStreamOS

> Distro differences

- Packaging
  - IPS: OI, Omnios, XStreamOS
  - SVR4: OpenSXCE, Tribblix
  - Deb: DilOS, osdyson
  - Pkgsrc: SmartOS
  - Storage vendors don't have (exposed) packaging
  - (No active rpm distro – Belenix is inactive)

> Distro differences

- X86/X64
  - Pretty much everything
- SPARC
  - OpenSXCE or Tribblix
- X64 only
  - SmartOS
- Minimum memory depends on packaging
  - IPS needs > 1G
  - Others happy in 512M

> Distro focus

- OpenIndiana – copy OpenSolaris
- OmniOS – server, supported
- OpenSXCE – copy Solaris 10
- XStreamOS – modern OpenIndiana
- SmartOS – foundation for cloud
- Tribblix – what OpenSolaris should have been

> Why do it yourself?

- Challenging and interesting
- Understand the inner workings
- Satisfy the target audience
- Make a flexible platform for development of new ideas
- Didn't like other distros!

> Tribblix values

- Modern components
- Retro styling
- Lightweight window managers
- SVR4 packaging
- Lightweight and fast
- Simplicity and “just works”

> Current state

- x86 variant pretty mature
- SPARC variant a work in progress
  - One of the few cross-platform distros
- Build process solidified
- Allowing further experimentation

> Release history

- Based around “milestones”
- Largely arbitrary placeholders
  - “Ooh, I got that working!”
- Milestone 0 October 2012
- Currently milestone 14
- So a release every ~2 months
  - Not quite as fast as I had hoped…

> Some Milestones

- Get it to boot at all!
- Build illumos; make self-hosting
- Functional desktop
- Gcc3 → gcc4 (needed 2 attempts)
- Network boot/install
- Openssl upgrade
- SPARC support

> One step at a time

- Started with OpenIndiana
- Convert into SVR4 packages
- Create bootable ISO image
  - That's a whole separate talk of itself
- Install by hand
  - Installer is cut and paste of the manual version
- Replace subsystems by our own
- SPARC started with OpenSXCE

> One-man show

- Is it feasible for one person to do this alone?
  - Essentially from scratch, not a distro clone
  - In their spare time, no less
- Initially estimated it would take 12 releases
  - Almost done at 14, had to back out and redo one
- Estimate of hours it would take was good
  - But actually finding time to do those hours was hard
- Took about twice as long as expected
  - Not that bad by industry standards!

> Create a live CD

- Most illumos distros are similar
  - Derived from Belenix
- Root archive as a ramdisk
- /usr as an iso image
- Scripting to glue them together
- Installer is a simple copy
  - Remove the glue
  - Add optional packages

> Integrated packaging

- Based on SVR4
  - Familiar so easy to use
  - Text files so trivial to manipulate
- Add wget/curl for downloads
  - And a catalog for inventory
- Aim is higher level administration
  - “Overlays” are clusters or groups
  - Packages are an implementation artefact

> Desktop options

- XFCE primary (now at 4.12)
- Enlightenment (was 0.17, now 0.19)
- Lots of others
  - WindowMaker, AfterStep, icewm, openbox, fluxbox
  - Twm, tvtwm, vtwm, piewm
  - Pekwm, amiwm, fvwm
  - CDE (partial port)
  - Awm (remember that?)

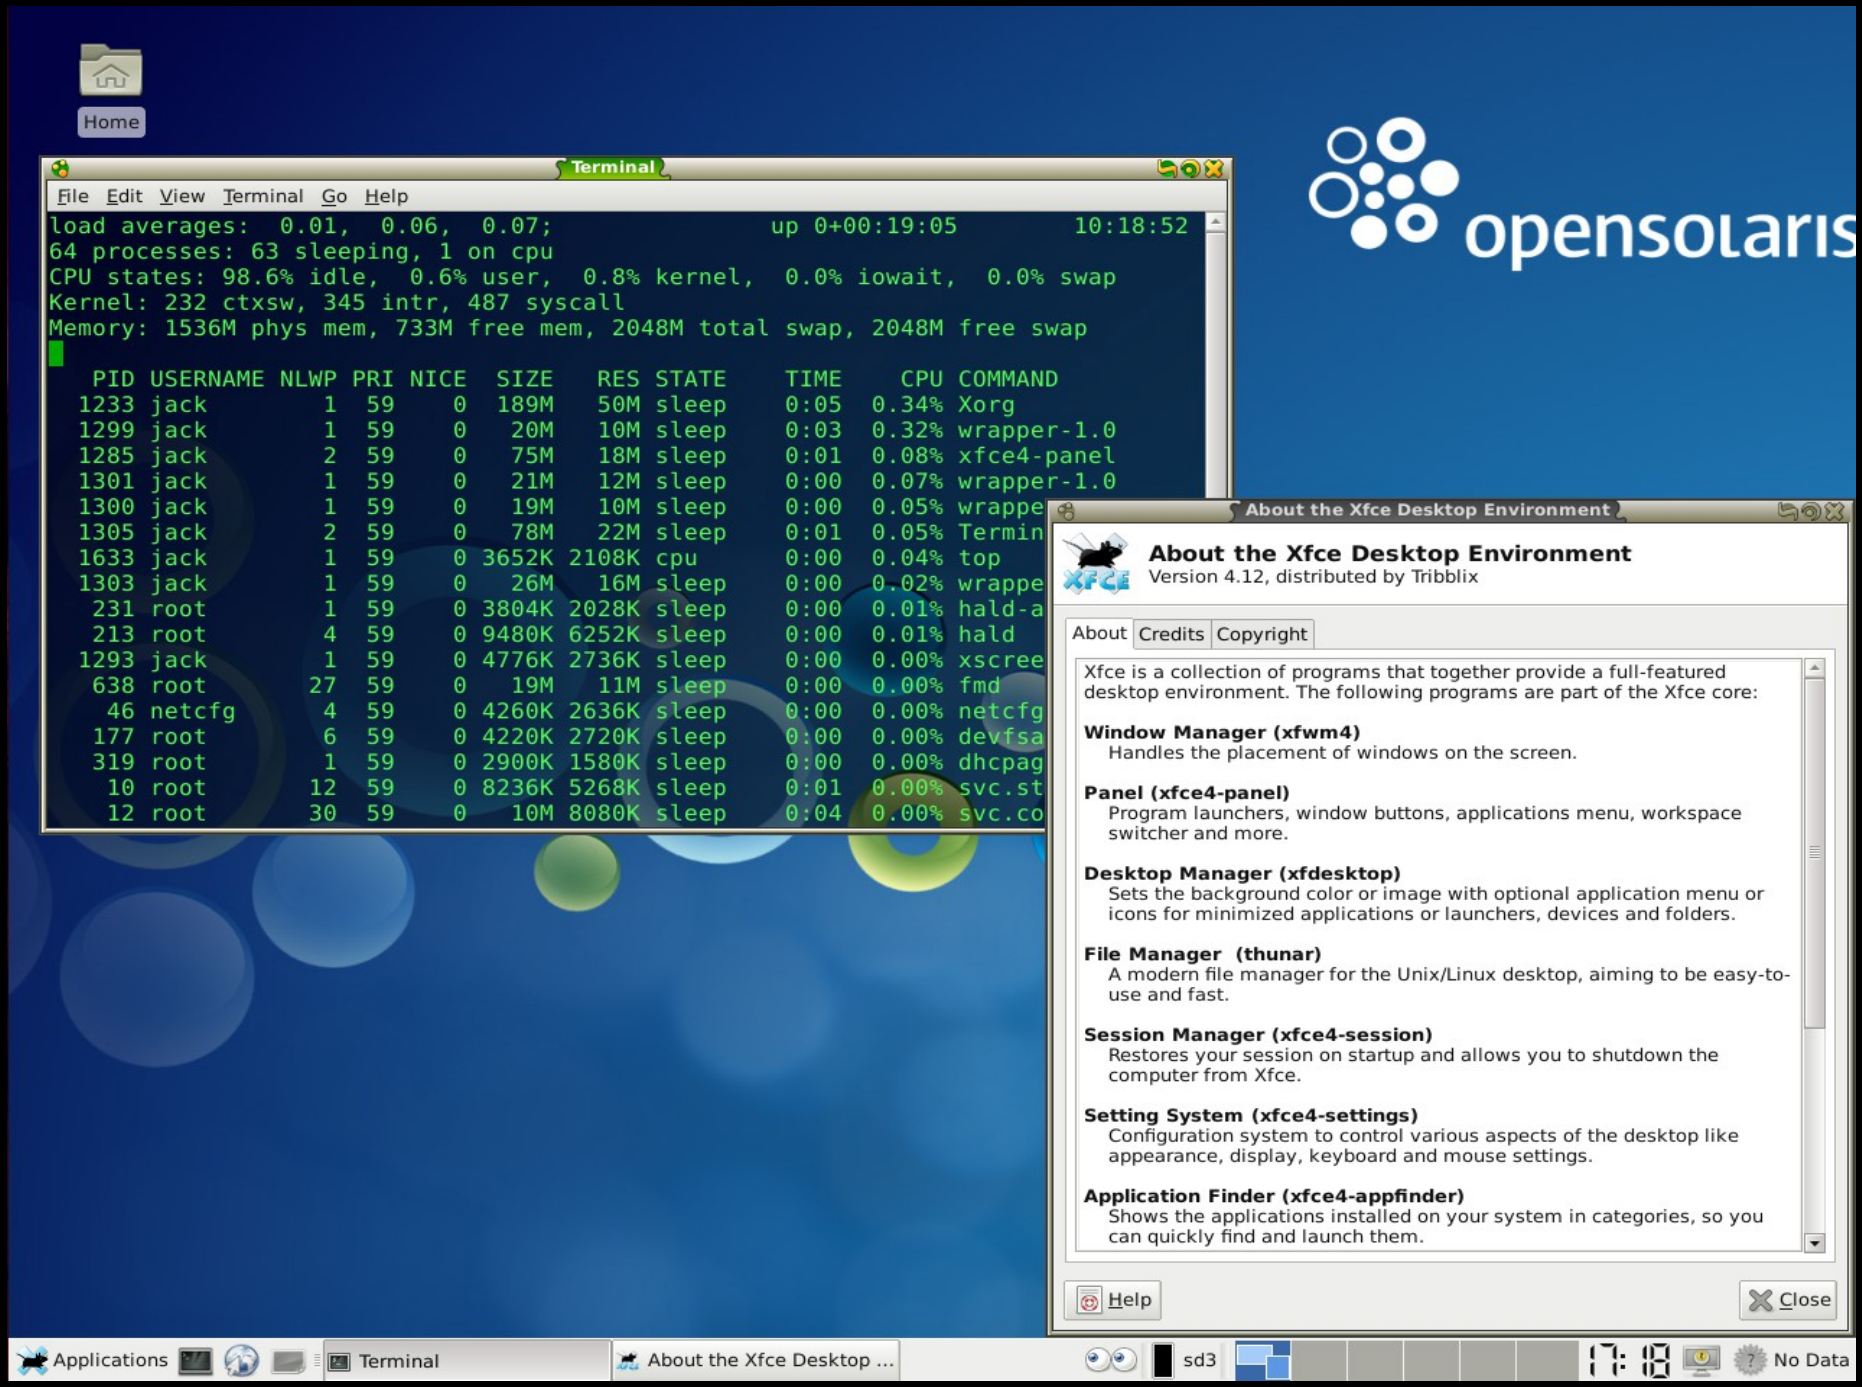
Home
Terminal
File Edit View Terminal Go Help
load averages: 0.01, 0.06, 0.07; up 0+00:19:05 10:18:52
64 processes: 63 sleeping, 1 on cpu
CPU states: 98.6% idle, 0.6% user, 0.8% kernel, 0.0% iowait, 0.0% swap
Kernel: 232 ctxsw, 345 intr, 487 syscall
Memory: 1536M phys mem, 733M free mem, 2048M total swap, 2048M free swap
PID USERNAME NLWP PRI NICE SIZE RES STATE TIME CPU COMMAND
1233 jack 1 59 0 189M 50M sleep 0:05 0.34% Xorg
1299 jack 1 59 0 20M 10M sleep 0:03 0.32% wrapper-1.0
1285 jack 2 59 0 75M 18M sleep 0:01 0.08% xfce4-panel
1301 jack 1 59 0 21M 12M sleep 0:00 0.07% wrapper-1.0
1300 jack 1 59 0 19M 10M sleep 0:00 0.05% wrappe
1305 jack 2 59 0 78M 22M sleep 0:01 0.05% Termin
1633 jack 1 59 0 3652K 2108K cpu 0:00 0.04% top
1303 jack 1 59 0 26M 16M sleep 0:00 0.02% wrappe
231 root 1 59 0 3804K 2028K sleep 0:00 0.01% hald-a
213 root 4 59 0 9480K 6252K sleep 0:00 0.01% hald
1293 jack 1 59 0 4776K 2736K sleep 0:00 0.00% xscree
638 root 27 59 0 19M 11M sleep 0:00 0.00% fmd
46 netcfg 4 59 0 4260K 2636K sleep 0:00 0.00% netcfg
177 root 6 59 0 4220K 2720K sleep 0:00 0.00% devfsa
319 root 1 59 0 2900K 1580K sleep 0:00 0.00% dhcpag
10 root 12 59 0 8236K 5268K sleep 0:01 0.00% svc.st
12 root 30 59 0 10M 8080K sleep 0:04 0.00% svc.co
opensolaris
About the Xfce Desktop Environment
About the Xfce Desktop Environment
Version 4.12, distributed by Tribblix
About Credits Copyright
Xfce is a collection of programs that together provide a full-featured desktop environment. The following programs are part of the Xfce core:
Window Manager (xfwm4)
Handles the placement of windows on the screen.
Panel (xfce4-panel)
Program launchers, window buttons, applications menu, workspace switcher and more.
Desktop Manager (xfdesktop)
Sets the background color or image with optional application menu or icons for minimized applications or launchers, devices and folders.
File Manager (thunar)
A modern file manager for the Unix/Linux desktop, aiming to be easy-to-use and fast.
Session Manager (xfce4-session)
Restores your session on startup and allows you to shutdown the computer from Xfce.
Setting System (xfce4-settings)
Configuration system to control various aspects of the desktop like appearance, display, keyboard and mouse settings.
Application Finder (xfce4-appfinder)
Shows the applications installed on your system in categories, so you can quickly find and launch them.
Help
Close
Applications
Terminal
About the Xfce Desktop ...
sd3
17:18
No Data


Home
Root
Temp
Terminal
File Edit View Terminal Go Help
Yes Peter? uname -a
SunOS tribblix 5.11 tribblix i86pc i386 i86pc
Yes Peter?
04:27 PM
Terminal

> Hypervisors all the way down

- Physical Domains (big iron)
- LDOM/VMware/Xen/KVM
- ipkg zones
- Whole root zones
- Containers/Docker/Rocket
- Sparse root zones
- Rump kernels

> The more things change...

- Solaris was lightweight
  - Moving up closer to a full VM
- Linux was heavyweight
  - Going all the way down with containers
- I've run Solaris Zones in production for a decade
- Zones could address the whole range
- Need to work on ease-of-use

> Zones in illumos

- Zones have 3 components
  - Kernel security/isolation
  - Commands for control
  - Brand – description and ancillary tools
- Illumos doesn't really provide any brands
- Distributions free to provide their own
  - IPS gives ipkg zones
  - SmartOS have a whole stack

> Zones in Tribblix

- Not implemented using packaging
  - Unlike Solaris 10 or IPS
- Reimplemented sparse root, whole root
  - Quick and simple
- Created partial root
  - Whole root with arbitrary software
- New alternate root
  - Sparse root with alternate origin
- Have a couple more types to come

Easy, right?

> Pain: building illumos

- Initially, always built on OpenIndiana
  - Still the default build environment for illumos
- Wanted to be self-hosting
- The illumos build is fragile
  - With specific dependencies
- Created a zone profile with all dependencies
  - Including IPS :-(
- Gives a reproducible controlled environment
  - Without polluting the rest of the system

> Pain: openssl

- Illumos has external dependencies
- OpenSSL is the thorniest
  - Used by SunSSH, sendmail, TPM, etc
- Much discussion about how to handle
- Official stance is to create a private copy
  - Which I've opposed
- Turns out the real problem is TPM
  - Which doesn't even come from illumos

> Pain: autotools

- Remember how autotools works...
  - Make a bunch of random guesses
  - Ignore the results
- Causes builds to be non-reproducible
- Need to explicitly direct configuration
  - The 'auto' is a misnomer
- Not yet satisfactorily solved
- Not that alternatives are any better

> Tribblix in flight

- Closing in on 1.0
  - Most initial milestones reached
- Bring SPARC to parity
- All core software Tribblix-native
  - Things like libxml2, libz – illumos depends on them
  - Much imported from OI or OpenSXCE
  - Takes time to build; often “improved” versions
- Key applications still need porting
  - LibreOffice, Firefox, GO

> Tribblix directions

- Zones and app deployment
  - Additional zone flavours in development
  - Fully integrated (think docker or app store)
- Simplify administration - “just works”
  - Make internals invisible
- Modern application stacks
  - Many on top of go
  - Integrated with zones and zfs

> illumos directions

- ZFS (OpenZFS)
- XPG7 standards
- Missing pieces from open code
- Cleaning cruft
  - But can we preserve heritage?
- LX brand (native Linux emulation)

> Potholes

- Not enough time/people
- Fragmentation
  - All the work done at the distro level
- SPARC port struggling
- No cgo yet

> Further reading

http://www.illumos.org/

http://www.tribblix.org/

http://www.petertribble.co.uk/

Peter.tribble@gmail.com